Dynamic control of sequential retrieval speed in networks with heterogeneous learning rules

Maxwell Gillett[1] and Nicolas Brunel[1,2*]

[1]Department of Neurobiology, Duke University, United States.
[2]Department of Physics, Duke University, United States.

*Corresponding author(s). E-mail(s): nicolas.brunel@duke.edu;

Temporal rescaling of sequential neural activity has been observed in multiple brain areas during behaviors involving time estimation and motor execution at variable speeds. Temporally asymmetric Hebbian rules have been used in network models to learn and retrieve sequential activity, with characteristics that are qualitatively consistent with experimental observations. However, in these models sequential activity is retrieved at a fixed speed. Here, we investigate the effects of a heterogeneity of plasticity rules on network dynamics. In a model in which neurons differ by the degree of temporal symmetry of their plasticity rule, we find that retrieval speed can be controlled by varying external inputs to the network. Neurons with temporally symmetric plasticity rules act as brakes and tend to slow down the dynamics, while neurons with temporally asymmetric rules act as accelerators of the dynamics. We also find that such networks can naturally generate separate 'preparatory' and 'execution' activity patterns with appropriate external inputs.

Introduction

Timing is a critical component in the proper planning and execution of temporally extended motor behaviors. In behaviors consisting of a single motor action, it may be desirable to control the duration of its execution. In behaviors composed of multiple actions, the precise time interval between actions can be a key determinant in the success of the behavior. How can the duration of these intervals be flexibly controlled in a network of neurons?

A simple mechanistic hypothesis for the regulation of motor related timing intervals posits a specialized neural circuit with network dynamics that vary

in speed as a consequence of differing levels of constant external input. Several network models utilizing external input as a means of speed control have been proposed to account for cortical and striatal dynamics observed during motor execution [1, 2]. To account for speed control in cortex, a recurrent neural network model has been trained to achieve temporal rescaling of network activity as a function of external input [2]. However, this model relies on supervised learning rules that may not be biologically plausible, and cannot generalize to other speeds from training on just one example timing interval. To explain speed control of sequential activity in striatum, a recurrent inhibitory network model has been proposed with a feedforward structure learned through anti-Hebbian plasticity [1]. This model demonstrates transient winner-take-all dynamics, with short term synaptic depression facilitating transitions in activity from one group of neurons to the next, and external input controlling the duration of each group's transient activation. While experimental evidence for the necessary type of depressive adaptation mechanism exists in the striatum, it may not be present in all cortical areas where rescaling of sequential activity is observed. Whether speed can be controlled in network models constructed using Hebbian learning without this mechanism remains unknown.

Network models with a connectivity generated by temporally asymmetric synaptic plasticity provide a potential framework for explaining how sequential activity can arise from local biologically plausible learning rules [3, 4]. In both rate and spiking networks, the temporal statistics of sequential activity in networks using this type of rule qualitatively match experimental findings made over both short and long timescales of observation in multiple tasks with timing components [5]. However, the speed of sequential dynamics in these models is constrained by the choice of temporal offset in the learning rule and neuronal time constant, and cannot be modulated with external input.

The Hebbian rules explored in this work and previous studies are approximations of various forms of spike-timing dependent plasticity (STDP). The effects of STDP can be quantified through kernels that measure the change in excitatory postsynaptic potential size at a synapse (as a proxy for synaptic strength), as a function of the timing difference between pre and postsynaptic spikes. Experimentally, a large diversity of STDP kernels have been characterized across cortical, subcortical, and cerebellar structures [6, 7]. Kernels measured in cortex and hippocampus typically, but not always, exhibit a temporal asymmetry, in which presynaptic activity must precede postsynaptic activity to elicit a positive change in synaptic strength [8, 9]. Theoretical studies have shown that this temporal asymmetry can be used to store and retrieve sequences of activity [1, 5, 10–20]. However, symmetric kernels, in which coincident activity leads to strengthening regardless of the order of pre and post-synaptic spikes, have also been observed in multiple contexts - with high frequency plasticity induction protocols in cortex [21], in hippocampal cultures in the presence of dopamine [22], and at excitatory-to-excitatory synapses in hippocampal CA3 [23]. Hebbian learning rules that are temporally symmetric lead instead to the creation of fixed point attractors [24–28].

It is not known to what degree temporal asymmetry varies across synapses at the scale of local networks, but analysis of a calcium-based plasticity model demonstrates that the degree of asymmetry can be controlled via adjustment of biophysical parameters [29]. We hypothesize that variability in the temporal offset expressed at a synapse may be a key ingredient in permitting the control of retrieval speed, suggesting a potential new role for the observed heterogeneity in STDP kernels.

In this work, we explore a generalization of previously investigated temporally asymmetric learning to multiple temporal offsets that captures this heterogeneity. Specifically, we find that varying the temporally asymmetry of the learning rule across synapses gives rise to network mechanisms that allow for the control of speed as a function of external inputs to the network. We start by considering a network with a bimodal distribution of heterogeneity in the learning rule, resulting in two distinct populations: one with a symmetric learning rule, and one with an asymmetric rule. We characterize the effect of input strength on retrieval speed and quality in these networks with connectivity generated using linear and nonlinear synaptic plasticity rules. We also find that transitions between fixed-point attractor-like "preparatory" periods and sequential "execution" phases can be realized in this model by rescaling the magnitude of external input. Finally, we demonstrate that networks with a uniform distribution of heterogeneity lead to qualitatively similar findings.

Degree of symmetry in learning rule determines retrieval speed

We explore a network model in which the firing rate dynamics of each neuron r_i in a population of size N is described by the equation

$$\tau \frac{dr_i}{dt} = -r_i + \phi\left(\sum_{j=1}^{N} J_{ij} r_j + I_i^{ext}(t)\right) \tag{1}$$

where τ is the time constant of firing rate dynamics, J_{ij} is the connectivity matrix, $\phi(x)$ is a sigmoidal neuronal transfer function (see Methods), and $I_i^{ext}(t)$ describes the external input provided to each neuron at time t.

We follow a similar learning procedure as in [5]. A sequence of P random i.i.d standard Gaussian patterns ξ_i^μ is presented to the network and stored in network connectivity. This sequence of patterns modifies the strength of synaptic connections J_{ij} from neuron j to i according to a Hebbian learning rule that transforms pre and post synaptic inputs into synaptic weight changes. The resulting connectivity matrix J_{ij} is a generalization of previously studied rules which combines both temporally symmetric and asymmetric learning [5, 28],

$$J_{ij} = A\frac{c_{ij}}{Nc}\left(\sum_{\mu=1}^{P} z_i f(\xi_i^\mu) g(\xi_j^\mu) + \sum_{\mu=1}^{P-1} (1-z_i) f(\xi_i^{\mu+1}) g(\xi_j^\mu)\right) \tag{2}$$

where c_{ij} is a matrix describing the structural connectivity, whose entries are given by i.i.d. Bernoulli random variables, $p(c_{ij} = 1) = c$, $p(c_{ij} = 0) = 1 - c$, where c is the connection probability; The functions $g(x)$ and $f(x)$ describe how the synaptic plasticity rule depends on pre and postsynaptic input patterns during learning, respectively; The parameter A controls the overall strength of the recurrent connections. And $z_i \in [0, 1]$ describes the degree of temporal symmetry at synapses of neuron i. A neuron with fully temporally symmetric plasticity is described by $z_i = 1$, while $z_i = 0$ indicates a neuron with fully temporally asymmetric plasticity. Note that we focus here to the case of a single sequence stored in synaptic connectivity, but such networks can also store multiple sequences [5].

We first explore the bilinear learning rule scenario ($f(x) = g(x) = x$) with homogeneous synaptic plasticity, i.e. $z_i = z$ for all $i = 1, \ldots, N$. At the two extremes of this variable we can recover previously studied learning rules. When $z = 0$, only the second term in Eq. 2 is present, resulting in a purely temporally asymmetric rule. Networks with connectivity constructed using such a rule can recall a sequence of stored patterns, and their sequential retrieval dynamics have been extensively characterized [5]. When $z = 1$, synaptic plasticity is temporally symmetric, potentially leading to fixed point attractor dynamics [28]. If z is instead fixed to a value between 0 and 1, then the asymmetric component in the plasticity rule leads to the retrieval of the whole sequence, but the speed at which the sequence is retrieved strongly depends on z. For instance, in Fig. 1b we demonstrate retrieval for an intermediate value of $z = 0.5$. Retrieval is quantified by plotting the Pearson correlation of the instantaneous firing rate $r(t)$ with each stored pattern ξ^μ as a function of time (see Methods). During sequence retrieval, correlations with individual patterns in the sequence increase, peak and decrease one after the other, indicating the network transiently visit states close to each of the patterns in succession. We find that in such a network, retrieval speed strongly depends on z. For the parameters in Fig. 1b, retrieval proceeds nearly twice as slowly as compared to a network with connectivity arising from a purely asymmetric learning rule, where retrieval speed is fixed by the time constant of the firing rate dynamics [5]. However, retrieval speed is fixed by the choice of z (see Fig. 1c showing a linear dependence of speed on z), and cannot be dynamically modulated in response to changes in the external input I_i^{ext}.

Heterogeneity in synaptic plasticity temporal asymmetry gives rise to a speed control mechanism

We next explored whether adding heterogeneity to this learning rule, allowing z_i to differ across synapses, can produce networks capable of both recalling stored sequences of patterns and modulating the speed of recall. We initially consider a bimodal distribution of degrees of temporal symmetry across the network. For each neuron, z_i was drawn randomly and independently as a

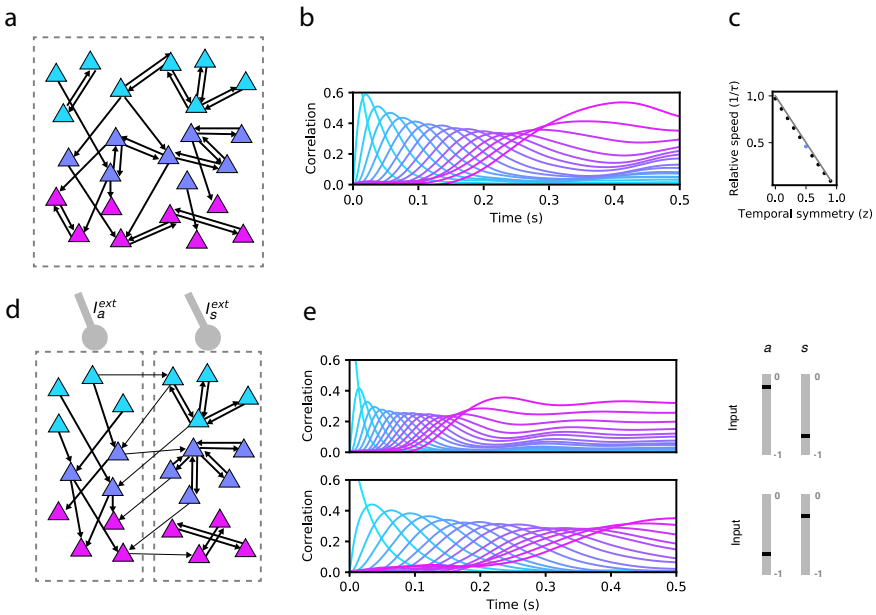

Fig. 1 Network model and sequence retrieval. **a**. Schematic of network connectivity after learning with a plasticity rule that combines temporally symmetric and asymmetric components. The network stores a sequence of patterns that activate non-overlapping sets of neurons (colored according to the pattern that activates them). Note connections both within each set, and from one set to the next. **b**. Correlation of each stored pattern with network activity following initialization to the first pattern. Retrieval speed is fixed by the balance of symmetry/asymmetry at the synapse. **c**. Relative retrieval speed as a function of temporal symmetry (z), showing linear relationship. Solid line: $1 - z$, the speed computed from MFT (see Methods). Black dots: Network simulations. **d**. Connectivity of a network with two types of neurons, asymmetric (left) and symmetric (right). Note that the connections from left neurons project to neurons active in the next pattern in the sequence, while connections from right neurons project to neurons active in the same pattern as the pre-synaptic neuron. The two types of neurons can be driven differentially by external inputs (I_a^{ext} and I_s^{ext}, respectively) **e**. Correlations as in (a) for two distinct pairs of input strengths (in the range [-1,0] for I_a^{ext} and I_s^{ext}), demonstrating two different retrieval speeds. In the simulations shown on the center and right panels, $N = 80{,}000$, $c = 0.005$, $\tau = 10$ms, $P = 16$, $A = 2$, $\theta = 0$, and $\sigma = 0.1$. For simplicity, we depict only 3 of the 16 stored patterns in the left schematics.

Bernoulli random variable with probability $p(z_i = 1) = 0.5$, $p(z_i = 0) = 0.5$. As a result, the network of N neurons can be divided into two subpopulations of approximately equal sizes $N_a = N_s = \frac{N}{2}$ neurons, according to the learning rule present at their synapses:

$$\tau \frac{dr_i^a}{dt} = -r_i^a + \phi\left(\sum_{j=1}^{N_a} J_{ij}^{aa} r_j^a + \sum_{j=1}^{N_s} J_{ij}^{as} r_j^s + I_a^{ext}(t)\right) \tag{3}$$

$$\tau \frac{dr_i^s}{dt} = -r_i^s + \phi\left(\sum_{j=1}^{N_s} J_{ij}^{ss} r_j^s + \sum_{j=1}^{N_a} J_{ij}^{sa} r_j^a + I_s^{ext}(t)\right) \tag{4}$$

where the connectivity matrix is given by

$$J_{ij}^{aX} = \frac{c_{ij}^{aX}}{N_a c} \sum_{\mu}^{P-1} f(\xi_i^{a,\mu+1}) g(\xi_j^{X,\mu}) \tag{5}$$

$$J_{ij}^{sX} = \frac{c_{ij}^{sX}}{N_s c} \sum_{\mu}^{P} f(\xi_i^{s,\mu}) g(\xi_j^{X,\mu}) \tag{6}$$

and where $X = a, s$ denotes the presynaptic population. Note that the external input I_X^{ext} now depends on the population. To reduce the space of possible learning rules, we have assumed that the type of learning at a synapse depends only on the identity of the postsynaptic neuron. The bimodal distribution of z_i restricts synapses to only one of the two types of plasticity, but in the final section entitled "Retrieval with a broad distribution of learning rules" we relax this constraint.

In Fig. 1e, we show an example of how the stored sequence can be retrieved under different input conditions. In both the top and bottom panels of 1e, network activity is initialized to the first pattern in the sequence, and a constant external input is provided to each subpopulation ("asymmetric" input I_a^{ext}, and "symmetric" input I_s^{ext}). In the top panel, the symmetric population is effectively silenced with strongly negative input, resulting in retrieval that lasts approximately τP, consistent with the dynamics being driven purely by the asymmetric component in the learning rule [5]. In the bottom panel, this input is no longer strongly negative, causing retrieval time to more than double, due to the effect of the symmetric population that tends to slow down the dynamics.

To characterize how retrieval time depends on these two sources of external input, we explored the space of the parameters defining the inputs to the network, I_a and I_s. In Fig. 2, we show the dependence of retrieval quality and speed on these variables. Retrieval quality is quantified by measuring the maximal correlation of the final pattern in the retrieved sequence. Retrieval speed is measured in units of the inverse of the neural time constant, τ^{-1}. It is computed by measuring the average inverse time between the peaks of consecutive correlations of the network state with consecutive patterns in the sequence. For example, a speed of 0.5 corresponds to an average time difference of 2τ between the peaks of the correlations of two consecutive retrieved patterns with network state. In the upper left quadrant of Fig. 2b, speed depends primarily on the strength of input to the symmetric population. Moving away from this region in the direction of increasing symmetric input, retrieval speed slows down to approximately 0.5. In the lower right quadrant, retrieval speed instead depends primarily on the strength of external input provided to the asymmetric population. As this negative input grows, retrieval speed becomes

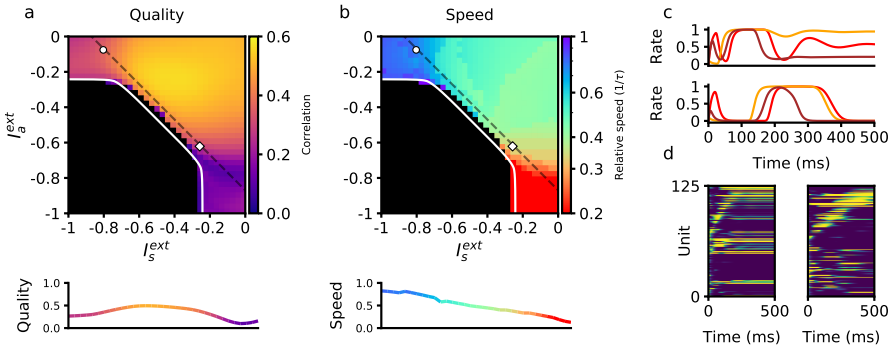

Fig. 2 Retrieval properties depend on external inputs. **a**. Retrieval quality, defined as the peak correlation m^P of the final pattern in the sequence, as a function of external inputs to asymmetric population I_a^{ext}, and to symmetric population I_s^{ext}. The white line bounds the region of successful retrieval. Below this line (black region), retrieval is not possible, regardless of initial condition (see Methods). **b**. Retrieval speed, measured by averaging the inverse time between consecutive pattern correlation peaks (see Methods). **c**. Firing rates of three randomly selected neurons during retrieval for parameters corresponding to the circle (left) and diamond (right) in panels (a-b). Note the approximate (but not exact) temporal scaling by a factor ~ 3 between these two sets of external inputs. **d**. Activity of 125 units (from a total of 80,000), sorted by peak firing rate time, for parameters corresponding to the circle (left) and diamond (right) in panels (a-b). All other parameters are as in Fig. 1b.

approximately four times slower than the speed of the purely asymmetric network. While we have focused on negative input in Fig. 2, retrieval speed is also modulated by positive input. Interestingly, it is the magnitude, not sign, of the input that determines retrieval speed. Expanding the phase diagram in panel (b) to positive input shows that the same dependence holds: values for retrieval speed are approximately symmetric about the I_a^{ext} and I_s^{ext} axes (not shown).

Flexible retrieval with a non-linear plasticity rule

We next considered the consequences of a nonlinear learning rule implemented by the following presynaptic and postsynaptic functions in Eq. 2:

$$f(x) = q_f - 1 + \Theta(x - x_f) \tag{7}$$

$$g(x) = q_g - 1 + \Theta(x - x_g) \tag{8}$$

where $\Theta(x)$ is the Heaviside function. This rule binarizes the activity patterns ξ according to a threshold, and its effects on persistent and sequential network activity have been studied extensively [5, 28, 30]. The parameter q_g is chosen such that $\int Dz g(z) = 0$, which keeps the mean connection strength at zero. The general dependency of retrieval speed on asymmetric and symmetric inputs in a network utilizing this rule is similar to that of the bilinear rule (see Fig. 2). One key difference is that a much wider range of speeds can be achieved using

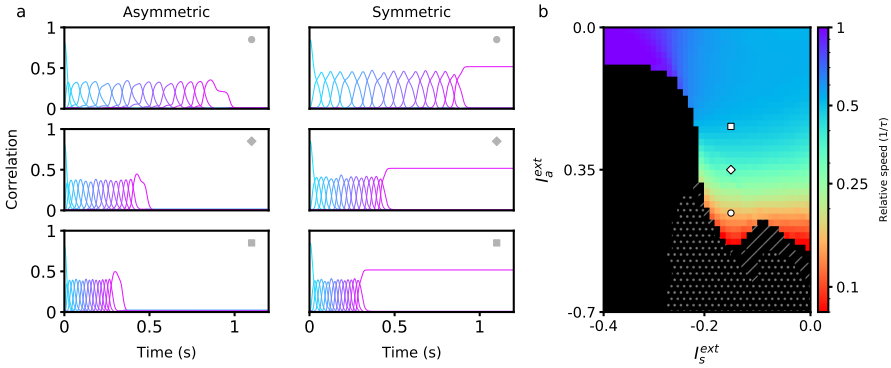

Fig. 3 Retrieval properties in networks with nonlinear learning rules. **a**. Correlations between stored patterns and network states in asymmetric (left) and symmetric (right populations, for three different external input combinations (denoted by the inset symbol, see right panel). **b**. Retrieval speed as a function of parameters describing external inputs, similarly as in Fig. 2. White dots indicate the region in which stable persistent activity of the first pattern is present. Hatched diagonal lines indicate the region in which incomplete sequential activity terminates in stable persistent activity. All parameters are as in Fig. 2, except $A = 20$ and $\sigma = 0.05$. The parameters of the learning rule are $x_f = 1.5$, $x_g = 1.5$, $q_f = 0.8$, and $q_g = 0.933$.

a nonlinear rule within the same retrieval quality bounds (see Methods). In fact, retrieval speed can now be arbitrarily slowed down, and even completely stopped when the input to the asymmetric population is sufficiently negative (see white dots in Fig. 3b). In this region, persistent activity is stable, and there exists a fixed point attractor correlated with any of the patterns in any stored sequence. There also exists a region in which sequential activity stops in the middle of retrieval and switches to stable persistent activity (see hatched diagonal lines in Fig. 3b).

Temporally varying external inputs can lead to transitions between persistent and sequential activity

We next explored how this heterogeneity might be used not only to control the speed of dynamics, but also to trigger transitions between qualitatively different dynamics. In Fig. 4 we use the same nonlinear model as in the previous section, and present discrete, time-dependent inputs intended to achieve persistent retrieval of a single pattern, followed by sequential retrieval of the remaining patterns at a specified time. To initiate persistent activity, we briefly present the first pattern as an input to the symmetric population. This elicits persistent activity in this population, as reflected by the sustained positive correlation of the symmetric population with the first pattern during the first 200ms (Fig. 4b). This activity does not recruit sequential activity in either

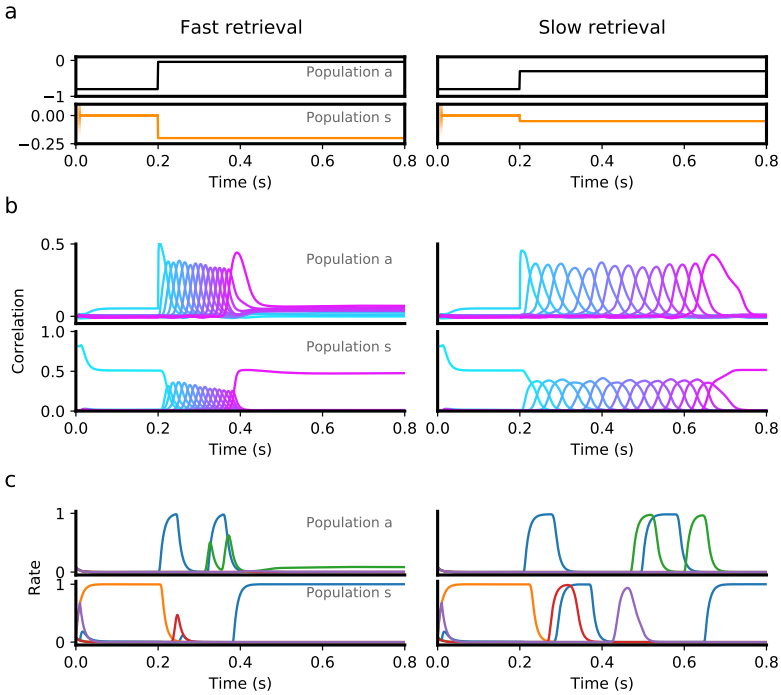

Fig. 4 Transition from persistent activity ('preparatory' period) to sequence retrieval ('execution' period) mediated by external input **a**. Inputs provided to the asymmetric (black) and symmetric population (orange) consist of a "preparatory period" input lasting 200ms, followed by an "execution period" input that is fixed for the rest of the interval. During a 200ms preparatory period, a brief input is presented to the symmetric population for the first 10ms, which drives the network to a state which is strongly correlated with the first pattern in a sequence. This input is removed after 10ms, but the network remains in a persistent activity state corresponding the the first pattern, because a strong negative input is presented to the asymmetric population throughout the entire 200ms, which prevents the network from retrieving the sequence. At the end of this period, the input to the symmetric population is decreased, while the asymmetric population is increased, which leads to retrieval of the sequence ('execution period'). Sequence retrieval can happen at different speeds, depending on the inputs to the asymmetric and symmetric populations. **b**. Correlations with stored patterns in the sequence in each population, in each input scenario. Note correlations in the slow retrieval case are temporally scaled by a factor ~ 2.5 compared to the fast retrieval case. **c**. Example single unit firing rates in each population. Note that for some neurons firing rates do not follow a simple temporal rescaling - for instance the purple neuron in the symmetric population is active at around $t = 0.45$ in the slow retrieval case, but is not active in the fast retrieval case. All parameters are as in Fig. 3, except $\theta = 0.07$ and $\sigma = 0.05$.

population, however, as the asymmetric population responsible for that transition is presented with sufficiently strong negative input during this period. To initiate sequential activity, inhibition to the asymmetric population is released after $t = 0.2$s, prompting the network to retrieve the stored sequence in both populations.

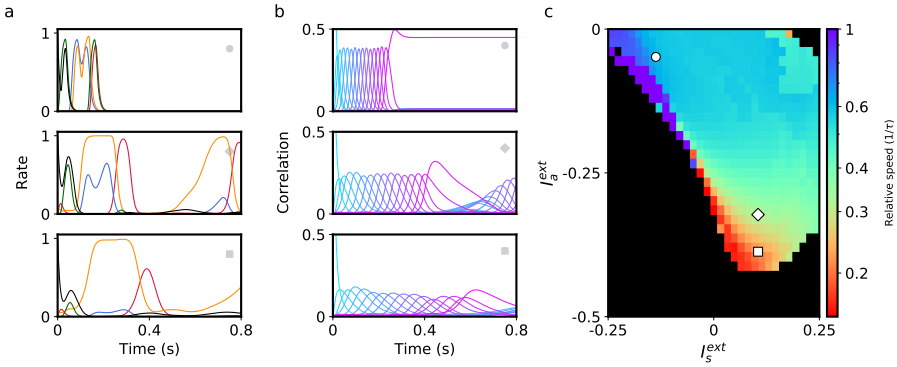

Fig. 5 Retrieval of sequences in networks with heterogeneous learning rules described by a continuum of degrees of symmetry. **a**. Firing rate dynamics of five representative neurons during retrieval for each external input configuration (see inset symbols in panel c). **b**. Correlation of network activity with each stored pattern during retrieval for each external input configuration. **c**. Retrieval speed as described in Fig. 2. All parameters are as in Fig. 3 except $A = 10$.

Note that in this scenario also, a sequence can be retrieved at various speeds, using the same inputs during the persistent period, but changing the level of constant stimulation provided during retrieval (compare left and right panels in Fig. 4b). As in a network with only a single asymmetric population, single neuron activity in this network is temporally sparse, with many neurons being active only at specific time intervals (Fig. 4c).

In our network, stability of persistent activity requires the dependence of the plasticity rule on pre and/or post synaptic firing rates to be non-linear. With a bilinear learning rule and Gaussian patterns, the network dynamics does not converge to fixed-point attractors that are correlated with a single pattern, but rather to mixed states correlated with multiple patterns [31].

The dynamics shown in Fig. 4 reproduces some of the landmark features observed in electrophysiological recordings during delayed motor tasks. In such tasks, a preparatory period follows presentation of a cue (e.g. instructing a target direction or a desired response speed), during which the animal can prepare the motor response, but not execute it [32]. This period is typically characterized by persistent activity of specific groups of neurons, whereas during motor execution those same neurons instead display transient activity [33].

Flexible sequence retrieval in networks with a continuous distribution of degrees of temporal symmetry

Up to this point, we have analyzed a network model in which neurons are separated in two discrete classes distinguished by their plasticity rule (symmetric or asymmetric). For a given postsynaptic neuron, the learning rule present at all

presynaptic synapses was chosen to be either temporally symmetric or asymmetric with equal probability, defining two distinct subpopulations of neurons. Can retrieval speed still be modulated by external input when synapses do not fall into such a binary classification, but have more heterogeneous properties? To model this heterogeneity, we chose to embed a continuum of learning rules. Instead of a bimodal distribution for z_i in Eq. 2, we choose a uniform distribution on the interval $[0, 1]$. The input I_i^{ext} provided to each neuron i in Eq. 1 is a linear combination of symmetric and asymmetric input components: $I_i^{ext} = z_i I_s^{ext} + (1 - z_i) I_a^{ext}$. We also choose to investigate a network with the previously described non-linear plasticity rule. Fig. 5 shows that a network with these modifications also exhibits flexible sequence retrieval, and that speed decreases as the asymmetric input component becomes more negative. However, as shown in Fig. 5c, to reach slower speeds a positive I_s^{ext} is now required. Note that a region of stable persistent activity is no longer present in this scenario, as stable persistent activity requires that a finite fraction of neurons in the network have a symmetric plasticity rule.

Discussion

In this paper, we have introduced a new mechanism for flexible control of retrieval speed in networks storing sequences. This mechanism relies on heterogeneity of synaptic plasticity rules across neurons in the network, with different degrees of temporal asymmetry. Neurons with temporally symmetric plasticity act as brakes of the dynamics, as they stabilize network activity in its current state, while neurons with temporally asymmetric plasticity act instead as accelerators, as they push the network towards the next pattern in the sequence. The speed of retrieval can then be modified in a flexible way by changing external inputs driving these two types of neurons. Furthermore, we found that this mechanism can be used to gate transitions between persistent and sequential activity.

Heterogeneity of synaptic plasticity

Our findings suggest a potential functional role for the experimentally observed diversity in synaptic plasticity rules [6–8, 21, 23]. In particular, a wide diversity of spike-timing dependent plasticity (STDP) curves have been reported in various brain structures, and sometimes in the same structure. In the hippocampus, temporally asymmetric STDP is typically observed in cultures [8] and at CA3 to CA1 connections in slices [34, 35], but temporally symmetric STDP is observed in area CA3 [23]. In the cerebellum, synaptic plasticity rules with diverse temporal requirements on the time difference between parallel fiber and climbing fiber inputs have been found in Purkinje cells in different zones of this structure [7]. While this heterogeneity has been found so far across structures or across different regions in the same structure, this heterogeneity could also be present within local networks, as current experimental methods for probing plasticity only have access to a single delay between pre

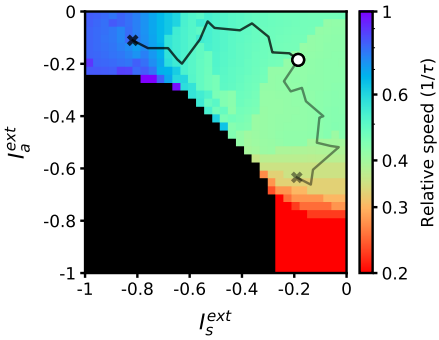

Fig. 6 Desired target speeds can be reached by adjusting external inputs using a reward-based learning rule. The black and grey lines denote the trajectories for two learning trials targeting different speeds. External inputs start at -0.2 (marked with a circle) and terminate at values implementing desired target speeds of 0.8 and 0.3 (marked with crosses). All parameters are as in Fig. 2.

and post-synaptic spikes in each recorded neuron, and would therefore miss this heterogeneity.

For simplicity, the degree of temporal asymmetry was chosen in our model to depend only on the identity of the postsynaptic neuron. This is consistent with the observation that a model of synaptic plasticity that depends only on the postsynaptic concentration of calcium can account for a range of experimentally observed STDP curves [29]. This suggests that heterogeneities in temporal asymmetry could arise due to heterogeneities in biophysical parameters that control calcium dynamics in post-synaptic spines.

Comparison with other mechanisms of speed control

The mechanism investigated here has major differences with previously described models of input-driven speed control. It does not require adaptation mechanisms or delays to slow down retrieval of subsequent patterns [1, 3]. It also does not require presentation of multiple exemplars spanning the desired range of retrieval speeds in order to find the appropriate network structure [2]. However, the mapping between external input strength and retrieval speed must be learned in order for the network to be able to perform retrieval at desired speeds. Unlike the model explored in [2], however, once this mapping is learned, it can be used to control the speed of other stored sequences.

Another recent study [36] has investigated how a recurrent network could flexibly control its temporal dynamics using a different approach. They trained a low-rank recurrent network using back-propagation through time to produce specific dynamics with flexible timing, and showed that the resulting network can then be flexibly controlled by a one-dimensional input. It would be interesting to investigate whether the low-rank structure found in such a manner exhibits similarities with the synaptic connectivity structure in our model.

The low-dimensional external inputs used to regulate speed are unrelated to the stored sequential input patterns, and suggest that a mapping to retrieval speed can be learned independently from a particular set of sequential patterns. We demonstrated that a reinforcement learning rule can be used to arrive at external input values implementing a desired speed (Fig. 6). By using a reward signal measuring how similar retrieval is to the desired speed, the rule adjusts initially random external inputs to the appropriate values over the course of multiple trial repetitions. Critically, once these external input values are learned, they can be used to modulate the retrieval speed of other stored sequences without having to relearn this mapping. Future experimental work could analyze the evolution of neural activity across the training of interval timing tasks, and evaluate whether it is consistent with such a reinforcement-based rule.

Experimental predictions

This mechanism presented here makes several predictions regarding the relationship between plasticity rules, external input, and the speed of network dynamics. One prediction is that retrieval speed could be modified by providing different external inputs to each population (asymmetric and symmetric). In vivo, these populations could be identified using the dependence of mean firing rates on speed of retrieval - neurons who increase their rates with slower/faster retrieval speeds would be predicted to be the symmetric/asymmetric neurons, respectively. Targeting one class of neurons or the other, using holographic techniques (see e.g. [37]) would then be expected to increase or decrease the speed of retrieval. Another prediction is that these cells have distinct profiles of temporal asymmetry in their synaptic plasticity. The model presented here also predicts the existence of "null" input directions, for which no change in retrieval speed is expected as external input is changed. When moving along these "null" directions, single neurons would only be expected to change their temporal firing patterns, but without affecting the speed of retrieval.

Transitions between persistent and sequential activity

Heterogeneity in the learning rule also provides a mechanism that enables input changes to drive transitions in activity states. An example of such a transition is frequently reported in primary motor cortex (M1) during delayed reaching tasks, where a preparatory period with persistent activity or ramping dynamics is followed by an execution period with transient, sequential dynamics [38, 39]. We demonstrated how an input change can gate such a transition in a simple network model composed of neurons with two distinct plasticity rules, the first temporally symmetric, and the second temporally asymmetric. At the start of the preparatory period, asymmetric neurons are inhibited, and a transient specific input elicits persistent activity in symmetric neurons. When inhibition is removed, asymmetric neurons become activated and drive a transition to sequential activity in both types of neurons.

Inhibitory gating has been previously hypothesized as a mechanism to control the initiation of execution period activity. Analysis of M1 activity suggests that local inhibitory interneurons do not engage in this gating, as putative inhibitory neurons do not appear to be preferentially active during the preparatory period compared to the execution period [40]. However, this does not rule out the possibility that the necessary inhibition could arise from other external inputs to M1. It is also possible that inhibition may not be required at all. Effective silencing of the asymmetric neurons could occur by a reduction of excitatory input during the preparatory period. Recent work in mice suggests that thalamocortical interactions may be a potential candidate for driving the required transition. Recorded activity in motor thalamus during a reaching task shows that at movement onset, thalamus activity is negatively correlated with premotor activity, but positively correlated with activity in M1 [41]. In a separate directional licking task, thalamus projections were shown to be required for initiating cued movement, and mimicked presentation of the cue when optogenetically stimulated [42]. An alternative model for transitions between preparatory and execution activity has recently been proposed [43], in which external inputs trigger a switch between a preparatory state and a nearly orthogonal execution state. However, in the model of ref. [43], the execution epoch is described by a single pattern, and any temporal dynamics within this epoch is inherited from external inputs.

Methods

Neuronal transfer function

The neuronal transfer function is given by the sigmoidal function

$$\phi(h) = \frac{r_{\max}}{2}\left(1 + \operatorname{erf}\left(\frac{h-\theta}{\sqrt{2}\sigma}\right)\right), \tag{9}$$

where θ determines the input at which the neuron fires at half the maximal value r_{max}, and σ is inversely proportional to the gain. This function was chosen for continuity with previous work [5]. We expect that using qualitatively similar functions should not alter the results of this paper.

Measuring pattern correlations

To compute the Pearson pattern correlation $m_\mu(t)$, we compute the overlap of each of the stored patterns ξ_μ with the instantaneous firing rates for the entire population and divide by the standard deviation of firing rate activity: $m_\mu(t) = \frac{1}{N}\sum_{i=1}^{N} r_i \xi_i^\mu / \sigma_r(t)$. In Figs. 3 and 4, we compute the correlations separately for each subpopulation.

Measuring retrieval speed

To measure retrieval speed v in Figs. 2, 3, and 5, we recorded the times at which each pattern correlation attained its peak value, and computed the average time difference between the peaks of successive patterns in a sequence. We then divided the time constant of the rate dynamics by this averaged value in order to convert speed into units of τ^{-1}:

$$v = \frac{\tau}{\frac{1}{P-1}\sum_{l=2}^{P} \operatorname{argmax}_t\big(m_l(t)\big) - \operatorname{argmax}_t\big(m_{l-1}(t)\big)} \tag{10}$$

To account for simulations with dynamics that did not have well-defined correlation peaks (typically observed at extreme storage loads or with persistent activity), we excluded peak time differences that exceeded two standard deviations of the average difference value. If no peak time difference passed this criteria, the sequence was considered not retrieved (black regions in Figs. 2, 3, and 5).

Mean-field theory of single-population network with variable degree of temporal asymmetry

In this section we derive a mean-field theory for the single population network with homogeneous synaptic plasticity. This is a generalization of the theory derived for a purely temporally asymmetric network [5]. We define order parameters $q_\mu(t) = \mathbb{E}\big(r(t)\xi^\mu\big)$ and $M = \mathbb{E}\big((r)^2(t)\big)$, describing the average overlap of network activity with pattern ξ^μ and the average squared firing rate, respectively.

Using equations (1,2), we derive equations describing the temporal evolution of the overlaps (for $l \in \{2, .., P\}$),

$$\tau\frac{dq_l}{dt} = -q_l + \int D\xi^l Dx\xi^l \phi\big((1-z)\xi^l q_{l-1} + z\xi^l q_l + R_l x + I^{\text{ext}}\big) \tag{11}$$

where R_l is a 'noise' term due to patterns $\mu \neq l$ in the sequence (see [5] for details) By making the following change of variables:

$$v = \frac{q_l^z \xi_l + xR_l}{\sqrt{(q_l^z)^2 + R_l^2}} \tag{12}$$

$$u = \frac{\xi_l R_l - q_l^z x}{\sqrt{(q_l^z)^2 + R_l^2}} \tag{13}$$

in which we have defined $q_l^z = (1-z)q_{l-1} + zq_l$, we obtain

$$\tau\frac{dq_l}{dt} = -q_l + q_l^z G\Big((q_l^z)^2 + R_l^2\Big), \tag{14}$$

where

$$G(x) = \frac{\int DuDv\phi\big(v\sqrt{x} + I_{ext}\big)}{\sqrt{x}}. \tag{15}$$

Assuming that $G(x) \approx 1$, which is the case during successful retrieval (see also [5]), then we can simplify to:

$$\frac{\tau}{1-z}\frac{dq_l}{dt} = -q_l + q_{l-1} \tag{16}$$

This equation makes it clear that retrieval speed depends linearly on z, i.e. on the balance between the symmetric and asymmetric components of synaptic plasticity.

Mean-field theory of heterogeneous network and conditions for retrieval

Mean-field theory can be used to further analyze retrieval speed dynamics, along the lines of [5]. We define order parameters $q_\mu^X(t) = \mathbb{E}\big(r^X(t)\xi^{X,\mu}\big)$ and $M_X = \mathbb{E}\big((r^X)^2(t)\big)$, describing the average overlap of network activity in subpopulation X with pattern $\xi^{X,\mu}$ and the average squared firing rate in subpopulation X, respectively. The equations for the overlaps are given by:

$$\tau\frac{dq_l^a}{dt} = -q_l^a + \big(q_{l-1}^a + q_{l-1}^s\big)\cdot G_a\Big((q_{l-1}^a + q_{l-1}^s)^2 + (R_l^a)^2\Big) \tag{17}$$

$$\tau\frac{dq_l^s}{dt} = -q_l^s + \big(q_l^s + q_l^a\big)\cdot G_s\Big((q_l^a + q_l^s)^2 + (R_l^s)^2\Big) \tag{18}$$

where G is given, for arbitrary transfer functions ϕ by:

$$G_X(x) = \frac{\int DuDv\phi\big(v\sqrt{x} + I_X^{ext}\big)}{\sqrt{x}}. \tag{19}$$

For the transfer function used in this paper, Eq. (9), the expression simplifies,

$$G_X(x) \equiv \frac{1}{\sqrt{2\pi(\sigma^2 + x)}}\exp\left(-\frac{(\theta + I_X^{ext})^2}{2(\sigma^2 + x)}\right). \tag{20}$$

As in the previous section, R_l^a and R_l^s are 'noise' terms due to patterns $\mu \neq l$ in the sequence, which also depends on the average squared firing rates M_a and M_s. Using Eqs. 17-18, we can derive the dynamics of the combined population overlap $q_l(t) = q_l^a(t) + q_l^s(t)$:

$$\tau\frac{dq^l}{dt} = -q^l + q_{l-1}G_a\Big((q_{l-1})^2 + (R_l^a)^2\Big) + q_l G_s\Big((q_l)^2 + (R_l^s)^2\Big) \tag{21}$$

To compute the boundary for successful retrieval given by the white line in Fig. 2, we analyze this equation when the gains are constant: $G_x(t) = G_x$. Plugging in and rearranging, we find:

$$\frac{\tau}{1-G_s}\frac{dq_l}{dt} = -q_l + \frac{G_a}{1-G_s}q_{l-1} \tag{22}$$

This equation shows that the sequence can only be retrieved if $G_a/(1-G_s) > 1$, otherwise the peak of the overlaps decay to zero with increasing l. Thus retrieval of an asymptotically long sequence is successful if the gain converges to a value greater or equal to one during retrieval. This condition can only be satisfied if

$$\max_x \left[G_a(x;\ \theta, \sigma, I^a_{ext}) + G_s(x;\ \theta, \sigma, I^s_{ext})\right] \geq 1 \tag{23}$$

To test for successful sequence retrieval in Fig. 2, we computed the maximal correlation value of the final pattern $m_P(t)$, and compared this value to a threshold $\theta_P = 0.05$. If the value fell below this threshold, then retrieval was considered unsuccessful, and was denoted by a black square. This threshold criterion was also used in Figures 3 and 5.

Reward-driven learning

A simple perturbation-based reinforcement learning rule is used to demonstrate that external inputs can be generated that produce network dynamics at a desired target speed over the course of multiple trial repetitions. We simulate a series of trials with stochastically varying external inputs. At each trial n, the external inputs used in the previous trial are perturbed randomly,

$$I^{ext,a}_{\text{pert}} = I^{ext,a}_{n-1} + \lambda \Delta x^a_n \tag{24}$$

$$I^{ext,s}_{\text{pert}} = I^{ext,s}_{n-1} + \lambda \Delta x^s_n \tag{25}$$

where λ is the strength of the perturbation, and Δx^p_n are uniformly distributed random variables over the interval $[-1, 1]$, drawn independently for each population $p \in \{a, s\}$ at each trial n. If these external inputs lead to an improvement in speed compared to previous trials, then

$$I^{ext,a}_n = I^{ext,a}_{\text{pert}} \tag{26}$$

$$I^{ext,s}_n = I^{ext,s}_{\text{pert}}; \tag{27}$$

else,

$$I^{ext,a}_n = I^{ext,a}_{n-1} \tag{28}$$

$$I^{ext,s}_n = I^{ext,s}_{n-1}. \tag{29}$$

In Fig. 6, the correlation threshold $\theta_m = 0.05$, the target speed $v_{target} = \{0.3, 0.8\}$, and $\lambda = 0.1$. On the first trial ($n = 0$), the external inputs are taken to be $I_0^{ext,a} = -0.2$ and $I_0^{ext,s} = -0.2$ (see open circle).

References

[1] Murray, J.M., Escola, G.S.: Learning multiple variable-speed sequences in striatum via cortical tutoring, 24 (2017)

[2] Wang, J., Narain, D., Hosseini, E.A., Jazayeri, M.: Flexible timing by temporal scaling of cortical responses. Nature Neuroscience **21**(1), 102–110 (2018). https://doi.org/10.1038/s41593-017-0028-6

[3] Sompolinsky, H., Kanter, I.: Temporal Association in Asymmetric Neural Networks. Physical Review Letters **57**(22), 2861–2864 (1986). https://doi.org/10.1103/PhysRevLett.57.2861

[4] Kleinfeld, D.: Sequential state generation by model neural networks. Proceedings of the National Academy of Sciences **83**(24), 9469–9473 (1986). https://doi.org/10.1073/pnas.83.24.9469

[5] Gillett, M., Pereira, U., Brunel, N.: Characteristics of sequential activity in networks with temporally asymmetric Hebbian learning. Proceedings of the National Academy of Sciences (2020). https://doi.org/10.1073/pnas.1918674117

[6] Abbott, L.F., Nelson, S.B.: Synaptic plasticity: Taming the beast. Nature Neuroscience **3**(11), 1178–1183 (2000). https://doi.org/10.1038/81453

[7] Suvrathan, A., Payne, H.L., Raymond, J.L.: Timing Rules for Synaptic Plasticity Matched to Behavioral Function. Neuron **92**(5), 959–967 (2016). https://doi.org/10.1016/j.neuron.2016.10.022

[8] Bi, G.-q., Poo, M.-m.: Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. Journal of Neuroscience **18**(24), 10464–10472 (1998). https://doi.org/10.1523/JNEUROSCI.18-24-10464.1998

[9] Egger, V., Feldmeyer, D., Sakmann, B.: Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. Nature Neuroscience **2**(12), 1098–1105 (1999). https://doi.org/10.1038/16026

[10] Jun, J.K., Jin, D.Z.: Development of Neural Circuitry for Precise Temporal Sequences through Spontaneous Activity, Axon Remodeling, and Synaptic Plasticity. PLoS ONE **2**(8), 723 (2007). https://doi.org/10.1371/journal.pone.0000723

[11] Liu, J.K., Buonomano, D.V.: Embedding Multiple Trajectories in Simulated Recurrent Neural Networks in a Self-Organizing Manner. Journal of Neuroscience **29**(42), 13172–13181 (2009). https://doi.org/10.1523/

JNEUROSCI.2358-09.2009

[12] Fiete, I.R., Senn, W., Wang, C.Z.H., Hahnloser, R.H.R.: Spike-Time-Dependent Plasticity and Heterosynaptic Competition Organize Networks to Produce Long Scale-Free Sequences of Neural Activity. Neuron **65**(4), 563–576 (2010). https://doi.org/10.1016/j.neuron.2010.02.003

[13] Waddington, A., Appleby, P.A., De Kamps, M., Cohen, N.: Triphasic spike-timing-dependent plasticity organizes networks to produce robust sequences of neural activity. Frontiers in Computational Neuroscience **6** (2012). https://doi.org/10.3389/fncom.2012.00088

[14] Zheng, P., Triesch, J.: Robust development of synfire chains from multiple plasticity mechanisms. Frontiers in Computational Neuroscience **8** (2014). https://doi.org/10.3389/fncom.2014.00066

[15] Okubo, T.S., Mackevicius, E.L., Payne, H.L., Lynch, G.F., Fee, M.S.: Growth and splitting of neural sequences in songbird vocal development. Nature **528**(7582), 352–357 (2015). https://doi.org/10.1038/nature15741

[16] Tannenbaum, N.R., Burak, Y.: Shaping neural circuits by high order synaptic interactions. PLOS Computational Biology **12**(8), 1005056 (2016). https://doi.org/10.1371/journal.pcbi.1005056

[17] Weissenberger, F., Meier, F., Lengler, J., Einarsson, H., Steger, A.: Long Synfire Chains Emerge by Spike-Timing Dependent Plasticity Modulated by Population Activity. International Journal of Neural Systems **27**(08), 1750044 (2017). https://doi.org/10.1142/S0129065717500447

[18] Theodoni, P., Rovira, B., Wang, Y., Roxin, A.: Theta-modulation drives the emergence of connectivity patterns underlying replay in a network model of place cells. eLife **7**, 37388 (2018). https://doi.org/10.7554/eLife.37388

[19] Pereira, U., Brunel, N.: Unsupervised Learning of Persistent and Sequential Activity. Frontiers in Computational Neuroscience **13** (2020). https://doi.org/10.3389/fncom.2019.00097

[20] Tupikov, Y., Jin, D.Z.: Addition of new neurons and the emergence of a local neural circuit for precise timing. Preprint, Neuroscience (March 2020)

[21] Sjöström, P.J., Turrigiano, G.G., Nelson, S.B.: Rate, Timing, and Cooperativity Jointly Determine Cortical Synaptic Plasticity. Neuron **32**(6), 1149–1164 (2001). https://doi.org/10.1016/S0896-6273(01)00542-6

[22] Zhang, J.C., Lau, P.M., Bi, G.Q.: Gain in sensitivity and loss in temporal

contrast of STDP by dopaminergic modulation at hippocampal synapses. Proc. Natl. Acad. Sci. U.S.A. **106**, 13028–13033 (2009)

[23] Mishra, R.K., Kim, S., Guzman, S.J., Jonas, P.: Symmetric spike timing-dependent plasticity at CA3–CA3 synapses optimizes storage and recall in autoassociative networks. Nature Communications **7**(1), 11552 (2016). https://doi.org/10.1038/ncomms11552

[24] Hopfield, J.J.: Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Sciences **79**(8), 2554–2558 (1982). https://doi.org/10.1073/pnas.79.8.2554

[25] Amit, D.J., Brunel, N.: Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cerebral Cortex **7**(3), 237–252 (1997). https://doi.org/10.1093/cercor/7.3.237

[26] Wang, X.-J.: Synaptic reverberation underlying mnemonic persistent activity. Trends in Neurosciences **24**(8), 455–463 (2001). https://doi.org/10.1016/S0166-2236(00)01868-3

[27] Brunel, N.: Course 10 - Network Models of Memory. In: Chow, C.C., Gutkin, B., Hansel, D., Meunier, C., Dalibard, J. (eds.) Les Houches. Methods and Models in Neurophysics, vol. 80, pp. 407–476. Elsevier, ??? (2005)

[28] Pereira, U., Brunel, N.: Attractor Dynamics in Networks with Learning Rules Inferred from In Vivo Data. Neuron **99**(1), 227–2384 (2018). https://doi.org/10.1016/j.neuron.2018.05.038

[29] Graupner, M., Brunel, N.: Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. Proceedings of the National Academy of Sciences **109**(10), 3991–3996 (2012). https://doi.org/10.1073/pnas.1109359109

[30] Lim, S., McKee, J.L., Woloszyn, L., Amit, Y., Freedman, D.J., Sheinberg, D.L., Brunel, N.: Inferring learning rules from distributions of firing rates in cortical neurons. Nature Neuroscience **18**(12), 1804–1810 (2015). https://doi.org/10.1038/nn.4158

[31] Amit, D., Gutfreund, H., Sompolinsky, H.: Spin-glass models of neural networks. Physical Review A **32**, 1007–1018 (1985)

[32] Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., Shenoy, K.V.: Neural population dynamics during reaching. Nature **487**(7405), 51–56 (2012). https://doi.org/10.1038/nature11129

[33] Svoboda, K., Li, N.: Neural mechanisms of movement planning: Motor cortex and beyond. Current Opinion in Neurobiology **49**, 33–41 (2018). https://doi.org/10.1016/j.conb.2017.10.023

[34] Wittenberg, G.M., Wang, S.S.: Malleability of spike-timing-dependent plasticity at the CA3-CA1 synapse. J. Neurosci. **26**, 6610–6617 (2006)

[35] Inglebert, Y., Aljadeff, J., Brunel, N., Debanne, D.: Synaptic plasticity rules with physiological calcium levels. Proc Natl Acad Sci U S A **117**, 33639–33648 (2020)

[36] Beiran, M., Meirhaeghe, N., Sohn, H., Jazayeri, M., Ostojic, S.: Parametric control of flexible timing through low-dimensional neural manifolds. Neuron **111**, 739–753 (2023)

[37] Marshel, J.H., Kim, Y.S., Machado, T.A., Quirin, S., Benson, B., Kadmon, J., Raja, C., Chibukhchyan, A., Ramakrishnan, C., Inoue, M., Shane, J.C., McKnight, D.J., Yoshizawa, S., Kato, H.E., Ganguli, S., Deisseroth, K.: Cortical layer-specific critical dynamics triggering perception. Science **365**(6453) (2019)

[38] Riehle, A., Requin, J.: Monkey primary motor and premotor cortex: Single-cell activity related to prior information about direction and extent of an intended movement. Journal of Neurophysiology **61**(3), 534–549 (1989). https://doi.org/10.1152/jn.1989.61.3.534

[39] Li, N., Daie, K., Svoboda, K., Druckmann, S.: Robust neuronal dynamics in premotor cortex during motor planning. Nature **532**(7600), 459–464 (2016). https://doi.org/10.1038/nature17643

[40] Kaufman, M.T., Churchland, M.M., Shenoy, K.V.: The roles of monkey M1 neuron classes in movement preparation and execution. Journal of Neurophysiology **110**(4), 817–825 (2013). https://doi.org/10.1152/jn.00892.2011

[41] Nashef, A., Mitelman, R., Harel, R., Joshua, M., Prut, Y.: Area-specific thalamocortical synchronization underlies the transition from motor planning to execution. Proceedings of the National Academy of Sciences of the United States of America **118**(6), 2012658118 (2021). https://doi.org/10.1073/pnas.2012658118

[42] Inagaki, H.K., Chen, S., Ridder, M.C., Sah, P., Li, N., Yang, Z., Hasanbegovic, H., Gao, Z., Gerfen, C.R., Svoboda, K.: A midbrain-thalamus-cortex circuit reorganizes cortical dynamics to initiate movement. Cell **185**(6), 1065–108123 (2022). https://doi.org/10.1016/j.cell.2022.02.006

[43] Bachschmid-Romano, L., Hatsopoulos, N., Brunel, N.: Interplay between

external inputs and recurrent dynamics during movement preparation and execution in a network model of motor cortex. eLife **in press**, 2022–0219481140